

# Shortest Paths in Intersection Graphs of Unit Disks

Sergio Cabello\*      Miha Jejčič†

21st February 2014

## Abstract

Let  $G$  be a unit disk graph in the plane defined by  $n$  disks whose positions are known. We show that in  $G$  a shortest path tree from a given source can be computed in  $\mathcal{O}(n \log n)$  time in the unweighted case and in  $\mathcal{O}(n^{1+\varepsilon})$  time in the weighted case. Previous algorithms used  $\mathcal{O}(n \log^6 n)$  expected time for the unweighted case and  $\mathcal{O}(n^{4/3+\varepsilon})$  time for the weighted case.

## 1 Introduction

Each set  $\mathcal{S}$  of geometric objects in the plane defines its intersection graph in a natural way: the vertex set is  $\mathcal{S}$  and there is an edge  $ss'$  in the graph,  $s, s' \in \mathcal{S}$ , whenever  $s \cap s' \neq \emptyset$ . It is natural to seek faster algorithms when the input is constraint to geometric intersection graphs. Here we are interested in computing shortest path distances in unit disk graphs, that is, the intersection graph of equal sized disks.

A unit disk graph is uniquely defined by the centers of the disks. Thus, we will drop the use of disks and just refer to the graph  $G(P)$  defined by a set  $P$  of  $n$  points in the plane. The vertex set of  $G(P)$  is  $P$ . Each edge of  $G(P)$  connects points  $p$  and  $p'$  from  $P$  whenever  $\|p - p'\| \leq 1$ , where  $\|\cdot\|$  denotes the Euclidean norm. See Figure 1 for an example of such graph. Up to a scaling factor,  $G(P)$  is isomorphic to a unit disk graph. In the *unweighted* case, each edge  $pp' \in E(G(P))$  has unit weight, while in the *weighted* case, the weight of each edge  $pp' \in E(G(P))$  is  $\|p - p'\|$ . In all our algorithms we assume that  $P$  is known. Thus, the input is  $P$ , as opposed to the abstract graph  $G(P)$ .

Exact computation of shortest paths in unit disks is considered by Roditty and Segal [15], under the name of *bounded leg* shortest path problem. They show that, for the weighted case, a shortest path tree can be computed in  $\mathcal{O}(n^{4/3+\varepsilon})$  time.

---

\*Department of Mathematics, IMFM, and Department of Mathematics, FMF, University of Ljubljana, Slovenia. Supported by the Slovenian Research Agency, program P1-0297, projects J1-4106 and L7-5459, and by the ESF EuroGIGA project (project GReGAS) of the European Science Foundation. E-mail: [sergio.cabello@fmf.uni-lj.si](mailto:sergio.cabello@fmf.uni-lj.si).

†Faculty of Mathematics and Physics, University of Ljubljana, Slovenia. E-mail: [jejicm@gmail.com](mailto:jejicm@gmail.com).

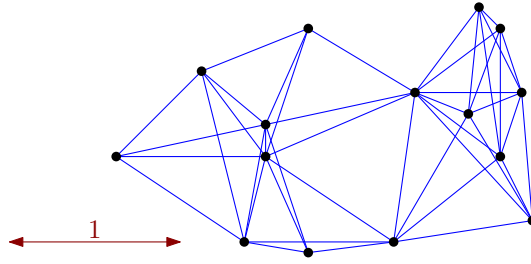


Figure 1: Example of graph  $G(P)$ .

They also note that the dynamic data structure for nearest neighbors of Chan [6] imply that, in the unweighted case, shortest paths can be computed in  $\mathcal{O}(n \log^6 n)$  expected time. (Roditty and Segal [15] also consider data structures to  $(1 + \varepsilon)$ -approximate shortest path distances in the intersection graph of congruent disks when the size of the disks is given at query time; they improve previous bounds of Bose et al. [4]. In this paper we do not consider that problem.)

We show that in  $G(P)$  a shortest path tree from a given source can be computed in  $\mathcal{O}(n \log n)$  time in the unweighted case and in  $\mathcal{O}(n^{1+\varepsilon})$  time in the weighted case. Thus, we provide a significant improvement over the results of Roditty and Segal. For the unweighted case we use the Delaunay triangulation to grow a search region and the algorithm is implementable. For the weighted case we use a modification of Dijkstra's algorithm combined with a data structure to dynamically maintain a bichromatic closest pair under an Euclidean weighted distance. Our algorithms are quite simple. Since the problem is very natural and we improve previous bounds, we think that simplicity is an added value.

Gao and Zhang [12] showed that the metric induced by a unit disk graph admits a compact well separated pair decomposition, extending the celebrated result of Callahan and Kosaraju [5] for Euclidean spaces. For making use of the well separated pair decomposition, Gao and Zhang [12] obtain a  $(1 + \varepsilon)$ -approximation to shortest path distance in unit disk graphs in  $\mathcal{O}(n \log n)$  time. Here we provide exact computation within comparable bounds.

Chan and Efrat [7] consider a graph defined on a point set but with more general weights in the edges. Namely, it is assumed that there is a function  $\ell: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}_+$  such that the edge  $pp'$  gets weight  $\ell(p, p')$ . Moreover, it is assumed that the function  $\ell(p, p')$  is increasing with  $\|p - p'\|$ . When  $\ell(p, p') = \|p - p'\|^2 f(\|p - p'\|)$  for a monotone increasing function  $f$ , then a shortest path can be computed in  $\mathcal{O}(n \log n)$  time. Otherwise, if  $\ell$  has constant size description, a shortest path can be computed in roughly  $\mathcal{O}(n^{4/3})$  time.

There has been a vast amount of work on algorithmic problems for unit disk and a review is beyond our possibilities. In the seminal paper of Clark, Colbourn and Johnson [8] it was shown that several  $\mathcal{NP}$ -hard optimization problems remain hard for unit disk graphs, although they showed the notable exception that maximum clique is solvable in polynomial time. Hochbaum and Maass [13] gave polynomial time approximation schemes for finding a largest independent set problems using the so-called shifting technique and there have been several developments since.

Shortest path trees can be computed for unit disk graphs in polynomial time.

One can just construct  $G(P)$  explicitly and run a standard algorithm for shortest paths. The main objective here is to obtain a faster algorithm that avoids the explicit construction of  $G(P)$  and exploits the geometry of  $P$ . There are several problems that can be solved in polynomial time, but faster algorithms are known for geometric settings. A classical example is the computation of the minimum spanning tree of a set of points in the Euclidean plane. Using the Delaunay triangulation, the number of relevant edges is reduced from quadratic to linear. For more advanced examples see Vaidya [16], Efrat, Itai and Katz [9], Eppstein [11], or Agarwal, Overmars and Sharir [3].

**Organization** In Section 2 we consider the unweighted case and in Section 3 we consider the weighted case. We conclude listing some open problems.

## 2 Unweighted shortest paths

In this section we consider the unweighted version of  $G(P)$  and compute a shortest path tree from a given point  $s \in P$ . Pseudocode for the eventual algorithm is provided in Figure 2. Before moving into the details, we provide the main ideas employed in the algorithm.

As it is usually done for shortest path algorithms we use tables  $\text{dist}[\cdot]$  and  $\pi[\cdot]$  indexed by the points of  $P$  to record, for each point  $p \in P$ , the distance  $d(s, p)$  and the ancestor of  $p$  in a shortest  $(s, p)$ -path. We start by computing the Delaunay triangulation  $DT(P)$  of  $P$ . We then proceed in rounds for increasing values of  $i$ , where at round  $i$  we find the set  $W_i$  of points at distance exactly  $i$  from the source  $s$ . We start with  $W_0 = \{s\}$ . At round  $i$ , we use  $DT(P)$  to grow a neighbourhood around the points of  $W_{i-1}$  that contains  $W_i$ . More precisely, we consider the points adjacent to  $W_{i-1}$  in  $DT(P)$  as candidate points for  $W_i$ . For each candidate point that is found to lie in  $W_i$ , we also take its adjacent vertices in  $DT(P)$  as new candidates to be included in  $W_i$ . For checking whether a candidate point  $p$  lies in  $W_i$  we use a data structure to find the nearest neighbour of  $p$  in  $W_{i-1}$ , denoted by  $\text{NN}(W_{i-1}, p)$ . Such data structure is just a point location data structure in the Voronoi diagram of  $W_{i-1}$ . Similarly, the shortest path tree is constructed by connecting each point of  $W_i$  to its nearest neighbour in  $W_{i-1}$ . See Figure 2 for the eventual algorithm UNWEIGHTEDSHORTESTPATH. In Figure 3 we show an example of what edges of the shortest path tree are computed in one iteration of the main loop.

We would like to emphasize a careful point that we employ to achieve the running time  $\mathcal{O}(n \log n)$ . In lines 16 and 17 of the algorithm, we check whether  $p$  is at distance at most 1 from *some* point in  $W_{i-1}$ , namely its nearest neighbour in  $W_{i-1}$ . Checking whether  $p$  is at distance at most 1 from  $\pi(q)$  (or  $q$  when  $q \in W_{i-1}$ ) would lead to a potentially larger running time. Thus, we do not grow each disk  $D(w, 1)$  independently for each  $w \in W_{i-1}$ , but we grow the whole region  $\bigcup_{w \in W_{i-1}} D(w, 1)$  at once. Growing each disk  $D(w, 1)$  separately would force us to check the same edge  $qp$  of  $DT(P)$  several times, once for each  $w \in W_{i-1}$  such that  $q \in D(w, 1)$ .

```

UNWEIGHTEDSHORTESTPATH( $P, s$ )
1  for  $p \in P$ 
2       $\text{dist}[p] = \infty$ 
3       $\pi[p] = \text{NIL}$ 
4   $\text{dist}[s] = 0$ 
5  build the Delaunay triangulation  $DT(P)$ 
6   $W_0 = \{s\}$ 
7   $i = 1$ 
8  while  $W_{i-1} \neq \emptyset$ 
9      build data structure for nearest neighbour queries in  $W_{i-1}$ 
10      $Q = W_{i-1}$     // candidate points
11      $W_i = \emptyset$ 
12     while  $Q \neq \emptyset$ 
13          $q$  an arbitrary point of  $Q$ 
14         remove  $q$  from  $Q$ 
15         for  $qp$  edge in  $DT(P)$ 
16              $w = \text{NN}(W_{i-1}, q)$ 
17             if  $\text{dist}[p] = \infty$  and  $\|p - w\| \leq 1$ 
18                  $\text{dist}[p] = i$ 
19                  $\pi[p] = w$ 
20                 add  $p$  to  $Q$ 
21                 add  $p$  to  $W_i$ 
22      $i = i + 1$ 
23 return  $\text{dist}[\cdot]$  and  $\pi[\cdot]$ 

```

Figure 2: Algorithm to compute a shortest path tree in the unweighted case.

**Lemma 1.** *Let  $p$  be a point from  $P \setminus \{s\}$  such that  $d(s, p) < \infty$ . There exists a point  $w$  in  $P$  and a path  $\pi$  in  $DT(P) \cap G(P)$  from  $w$  to  $p$  such that  $d(s, w) + 1 = d(s, p)$  and each internal vertex  $p_j$  of  $\pi$  satisfies  $d(s, p_j) = d(s, p)$ .*

*Proof.* Let us set  $i = d(s, p)$ . Let  $w$  be the point with  $d(s, w) = i - 1$  that is closest to  $p$  in Euclidean distance. It must be that  $\|w - p\| \leq 1$  because  $d(s, p) < \infty$ . Let  $D_{wp}$  be the disk with diameter  $wp$ .

For simplicity, let us assume that the segment  $wp$  does not go through any vertex of the Voronoi diagram of  $P$ . (In the degenerate case where this does not happen, we can replace  $p$  by a point  $p'$  arbitrarily close to  $p$ .) Consider the sequence of Voronoi cells  $\text{cell}(p_1, P), \dots, \text{cell}(p_k, P)$  intersected by the segment  $wp$ , as we walk from  $w$  to  $p$ . See Figure 4. Clearly  $w = p_1$  and  $p = p_k$ . For each  $1 \leq j < k$ , the edge  $p_j p_{j+1}$  is in  $DT(P)$  because  $\text{cell}(p_j, P)$  and  $\text{cell}(p_{j+1}, P)$  are adjacent along some point of  $wp$ . Therefore the path  $\pi = p_1 p_2 \dots p_k$  is contained in  $DT(P)$  and connects  $w$  to  $p$ . For any index  $j$  with  $1 < j < k$ , let  $a_j$  be any point in  $wp \cap \text{cell}(p_j, P)$ . Since  $\|a_j p_j\| \leq \min\{\|a_j w\|, \|a_j p\|\}$ , the point  $p_j$  is contained in  $D_{uw}$ . Therefore the whole path  $\pi$  is contained in  $D_{wp}$  and, since  $D_{wp}$  has diameter at most one, each edge of  $\pi$  is also in  $G(P)$ . We conclude that  $\pi$  is a path in

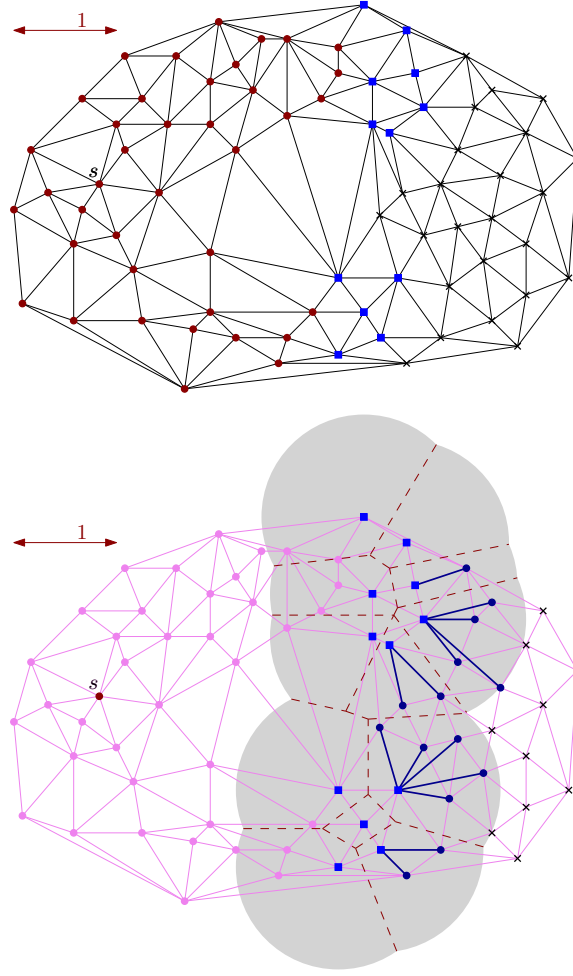


Figure 3: Top: A point set with its Delaunay triangulation. The source is marked as  $s$ . Points  $p$  such that  $d(s, p) \leq 3$  are marked with red dots. Points from  $W_4$  are marked with blue boxes. Bottom: The new edges added to the tree at iteration 5 and the new vertices are shown. The light grey region is  $\bigcup_{p \in W_4} D(p, 1)$ , where the Voronoi diagram of  $W_4$  is superimposed.

$DT(P) \cap G(P)$ .

Consider any point  $p_j$  of  $\pi$ , which is thus contained in  $D_{wp}$ . Because  $\|w - p_j\| \leq \|w - p\| \leq 1$ , we have  $d(s, p_j) \leq d(s, w) + 1 = i$ . Because  $\|p_j - p\| \leq \|w - p\| \leq 1$ , we have  $d(s, p_j) \geq d(s, p) - 1 = i - 1$ . However, the choice of  $w$  as closest to  $p$  implies that  $d(s, p_j) \neq i - 1$  because  $\|p_j - p\| < \|w - p\|$ . Therefore  $d(s, p_j) = i$ . We conclude that all internal vertices  $p_j$  of  $\pi$  satisfy  $d(s, p_j) = i$ .  $\square$

**Lemma 2.** *At the end of algorithm UNWEIGHTEDSHORTESTPATH( $P, S$ ) it holds*

$$\forall i \in \mathbb{N} \cup \{0\} : \quad W_i = \{p \in P \mid d(s, p) = i\}.$$

*Moreover, for each point  $p \in P \setminus \{s\}$ , it holds that  $\text{dist}[p] = d(s, p)$  and, if  $d(s, p) < \infty$ , there is a shortest path in  $G(P)$  from  $s$  to  $p$  that uses  $\pi[p]p$  as last edge.*

*Proof.* We proof the statement by induction on  $i$ .  $W_0 = \{s\}$  is set in line 6 and never changed. Thus the statement holds for  $i = 0$ .

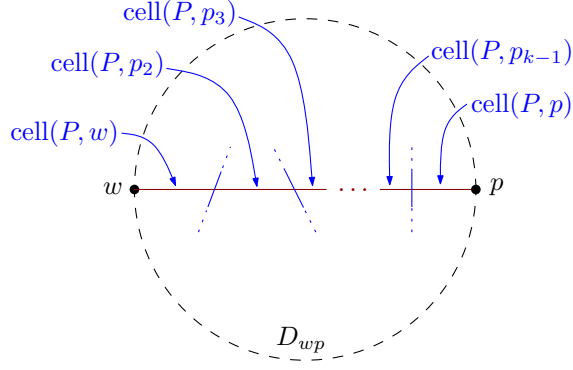


Figure 4: Proof of Lemma 1.

Before considering the inductive step, note that the sets  $W_0, W_1, \dots$  are pairwise disjoint. Indeed, a point  $p$  is added to some  $W_i$  (line 21) at the same time that we set  $\text{dist}[p] = i$  (line 18). After setting  $\text{dist}[p]$ , the test in line 17 is always false and  $p$  is not added to any other set  $W_j$ .

Consider any value  $i > 0$ . By induction we have that

$$W_{i-1} = \{p \in P \mid d(s, p) = i - 1\}.$$

In the algorithm we add points to  $W_i$  only in line 21. If a point  $p$  is added to  $W_i$ , then  $\|p - w\| \leq 1$  for some  $w \in W_{i-1}$  because of the test in line 17. Therefore any point  $p$  added to  $W_i$  satisfies  $d(s, p) \leq i$ . Since  $p \notin W_{i-1}$ , the disjointness of the sets  $W_0, W_1, \dots$ , implies that  $d(s, p) = i$ . We conclude that

$$W_i \subseteq \{p \in P \mid d(s, p) = i\}.$$

For the reverse containment, let  $p$  be any point such that  $d(s, p) = i$ . We have to show that  $p$  is added to  $W_i$  by the algorithm. Consider the point  $w$  and the path  $\pi = p_1 \dots p_k$  guaranteed by Lemma 1. By the induction hypothesis,  $w = p_1 \in W_{i-1}$  and thus is added to  $Q$  in line 10. At some moment the edge  $p_1 p_2$  is considered in line 15 and the point  $p_2$  is added to  $W_i$  and  $Q$ . An inductive argument thus shows that all the points  $p_3, \dots, p_k$  are added to  $W_i$  and  $Q$  (possibly in a different order). It follows that  $p_k = p$  is added to  $W_i$  and thus

$$W_i = \{p \in P \mid d(s, p) = i\}.$$

Since a point  $p$  is added to  $W_i$  at the same time that  $\text{dist}[p] = i$  is set, it follows that  $\text{dist}[p] = i = d(s, p)$ . Since  $\pi[p] \in W_{i-1}$  and  $\|p - \pi[p]\| \leq 1$  (lines 16, 17 and 19), there is a shortest path in  $G(P)$  from  $s$  to  $p$  that uses an  $(i-1)$ -edge path from  $s$  to  $\pi[p]$ , by induction, followed by the edge  $\pi[p]p$ .  $\square$

**Lemma 3.** *The algorithm UNWEIGHTEDSHORTESTPATH( $P, S$ ) takes  $\mathcal{O}(n \log n)$  time, where  $n$  is the size of  $P$ .*

*Proof.* The main observations used in the proof are the following: each point of  $P$  is added to  $Q$  at most once in line 10 and once in line 20, the execution of lines 13–21 for a point  $q$  takes time  $\mathcal{O}(\deg_{DT(P)}(q) \log n)$ , the sum of the degrees of the

points in  $DT(P)$  is  $\mathcal{O}(n)$ , and in line 9 we spend time  $\mathcal{O}(n \log n)$  overall iterations together. We next provide the details.

The Delaunay triangulation of  $n$  points can be computed in  $\mathcal{O}(n \log n)$  time. Thus the initialization in lines 1–7 takes  $\mathcal{O}(n \log n)$  time. It remains to argue that the loop in lines 8–22 takes time  $\mathcal{O}(n \log n)$ .

An execution of the lines 9–11 takes time  $\mathcal{O}(|W_{i-1}| \log |W_{i-1}|) = \mathcal{O}(|W_{i-1}| \log n)$ . Each subsequent nearest neighbour query takes  $\mathcal{O}(\log n)$  time.

Each execution of the lines 16–21 takes time  $\mathcal{O}(\log n)$ , where the most demanding step is the query made in line 16. Each execution of the lines 13–21 takes time  $\mathcal{O}(\deg_{DT(P)}(q) \cdot \log n)$  because the lines 16–21 are executed  $\deg_{DT(P)}(q)$  times.

Consider one execution of the lines 9–22 of the algorithm. Points are added to  $Q$  in lines 10 and 20. In the latter case, a point  $p$  is added to  $Q$  if and only if it is added to  $W_i$  (line 21). It follows that a point is added to  $Q$  if and only if it belongs to  $W_{i-1} \cup W_i$ . Moreover, each point of  $W_{i-1} \cup W_i$  is added exactly once to  $Q$ : each point  $p$  that is added to  $Q$  has  $\text{dist}[p] \leq i < \infty$  and will never be added again because of the test in line 17. It follows that the loop in lines 12–22 takes time

$$\sum_{q \in W_{i-1} \cup W_i} \mathcal{O}(\deg_{DT(P)}(q) \cdot \log n).$$

Therefore we can bound the time spent in the the loop of lines 8–22 by

$$\sum_i \mathcal{O} \left( |W_i| \log n + \sum_{q \in W_{i-1} \cup W_i} (\deg_{DT(P)}(q) \cdot \log n) \right). \quad (1)$$

Using that the sets  $W_0, W_1, \dots$  are pairwise disjoint (Lemma 2) with  $\sum_i |W_i| \leq n$  and

$$\sum_{q \in P} \deg_{DT(P)}(q) = 2 \cdot |E(DT(P))| = \mathcal{O}(n),$$

the bound in (1) becomes  $\mathcal{O}(n \log n)$ .  $\square$

**Theorem 4.** *Let  $P$  be a set of  $n$  points in the plane and let  $s$  be a point from  $P$ . In time  $\mathcal{O}(n \log n)$  we can compute a shortest path tree from  $s$  in the unweighted graph  $G(P)$ .*

*Proof.* Consider the algorithm UNWEIGHTEDSHORTESTPATH( $P, S$ ) given in Figure 2. Because of Lemma 3 it takes time  $\mathcal{O}(n \log n)$ . Because of Lemma 2, the table  $\pi[\cdot]$  correctly describes a shortest path tree from  $s$  in  $G(P)$  and  $\text{dist}[\cdot]$  correctly describes shortest path distances in  $G(P)$ .  $\square$

### 3 Weighted shortest paths

In this section we consider the SSSP problem on the weighted version of  $G(P)$ : points  $p$  and  $q$  have an edge between them iff  $\|p - q\| \leq 1$  and the weight of that edge is  $\|p - q\|$ . Our algorithm uses a dynamic data structure for bichromatic closest pairs. We first review the precise data structure that we will employ. We then describe the algorithm and discuss its properties.

### 3.1 Bichromatic closest pair

Let  $P$  be a set of  $n$  points in the plane and let each point  $p \in P$  have a weight  $w_p \geq 0$ . We call a function  $\delta: \mathbb{R}^2 \times P \rightarrow \mathbb{R}_+$  a *(additive) weighted Euclidean metric*, if it is of the form

$$\delta(q, p) = w_p + \|q - p\|,$$

where  $\|\cdot\|$  denotes the Euclidean distance.

Let  $\varepsilon > 0$  denote an arbitrary constant. Agarwal, Efrat and Sharir [2] showed that for any  $P$  and  $\delta$  as above,  $P$  can be preprocessed in  $\mathcal{O}(n^{1+\varepsilon})$  time into a data structure of size  $\mathcal{O}(n^{1+\varepsilon})$  so that points can be inserted into or deleted from  $P$  in  $\mathcal{O}(n^\varepsilon)$  amortized time per update, and a nearest-neighbour query can be answered in  $\mathcal{O}(\log n)$  time. Eppstein [10] had already shown that if such a dynamic data structure existed, then a bichromatic closest pair (BCP) under  $\delta$  of red and blue points in the plane could be maintained, adding only a polylogarithmic factor to the update time. Combining these two results gives

**Theorem 5** (Agarwal, Efrat, Sharir [2]). *Let  $R$  and  $B$  be two sets of points in the plane with a total of  $n$  points. We can store  $R \cup B$  in a dynamic data structure of size  $\mathcal{O}(n^{1+\varepsilon})$  that maintains a bichromatic closest pair in  $R \times B$ , under any weighted Euclidean metric, in  $\mathcal{O}(n^\varepsilon)$  amortized time per insertion or deletion.*  $\square$

### 3.2 Algorithm

We will use a variant of Dijkstra’s algorithm. As before, we maintain tables  $\text{dist}[\cdot]$  and  $\pi[\cdot]$  containing distances from the source and parents of points in the shortest path tree. As in Dijkstra’s algorithm we will maintain a set  $S$  (containing the source  $s$ ) of points for which the correct distance from  $s$  has already been computed, and a set  $P \setminus S$  of points for which the distance has yet to be computed. For the points of  $S$ ,  $\text{dist}[\cdot]$  stores the true distance from the source.

In our approach we split  $S$  into sets  $B$  and  $D$ , called “blue” and “dead” points, respectively. We call the points in  $R = P \setminus S$  “red” points. The reason for the introduction of the “dead” points  $D$  is that, as it will be proved, at no stage of the algorithm there is an edge of  $G(P)$  between a point in  $D$  and a point in  $R$ . Thus, the points of  $D$  are not relevant to find the last edge in a shortest path to points of  $R$ .

We store  $R \cup B$  in the dynamic data structure from Theorem 5 that maintains the bichromatic closest pair (BCP) in  $R \times B$  under the weighted Euclidean metric

$$\delta(r, b) := \text{dist}[b] + \|r - b\|.$$

At each iteration of the main **while** loop, we query the data structure for a BCP pair  $(b^*, r^*)$ . If  $b^*r^*$  is not an edge in our underlying graph  $G(P)$ , meaning  $\|b^* - r^*\| > 1$ , then  $b^*$  will never be the last vertex to any point in  $R$ , and therefore we will move it from  $B$  to  $D$ . If  $b^*r^*$  is an edge of  $G(P)$ , then, as it happens with Dijkstra’s algorithm, we have completed a shortest path to  $r^*$ . The algorithm is given in Figure 5. Figure 6 shows sets  $D$ ,  $B$ , and  $R$  in the middle of a run of the algorithm.



WEIGHTEDSHORTESTPATHS( $P, s$ )

```

1  for  $p \in P$ 
2       $\text{dist}[p] = \infty$ 
3       $\pi[p] = \text{NIL}$ 
4   $\text{dist}[s] = 0$ 
5   $B = \{s\}$ 
6   $D = \emptyset$ 
7   $R = P \setminus \{s\}$ 
8  store  $R \cup B$  in the BCP dynamic DS of Theorem 5 wrt  $\delta(r, b)$ 
9  while  $R \neq \emptyset$ 
10     if  $B = \emptyset$ 
11         return  $\text{dist}[\cdot]$  and  $\pi[\cdot]$            //  $G(P)$  is not connected
12     else  $(b^*, r^*) = \text{BCP}(B, R)$ 
13     if  $\|b^* - r^*\| > 1$ 
14         delete( $B, b^*$ )
15          $D = D \cup \{b^*\}$ 
16     else  $\text{dist}[r^*] = \text{dist}[b^*] + \|b^* - r^*\|$ 
17          $\pi[r^*] = b^*$ 
18         delete( $R, r^*$ )
19         insert( $B, r^*$ )
20 return  $\text{dist}[\cdot]$  and  $\pi[\cdot]$ 

```

Figure 5: Algorithm for SSSP in the weighted case.

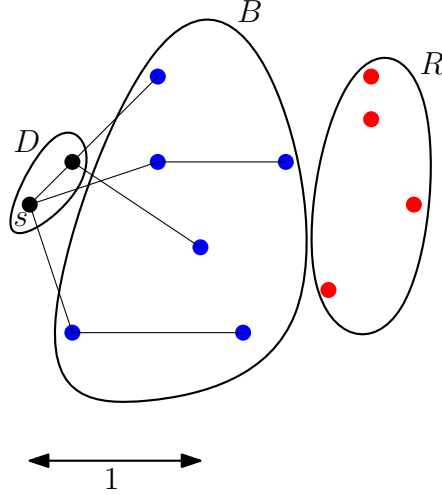


Figure 6: Sets  $D$ ,  $B$ , and  $R$  after a couple of iterations of the **while** loop in WEIGHTEDSHORTESTPATHS.

Let us explain the actual bottleneck of our approach to reduce the time from  $\mathcal{O}(n^{1+\epsilon})$  to  $\mathcal{O}(n \text{ polylog } n)$ . The inner workings of the data structure of Theorem 5 is

based on two dynamic data structures. One of them has to compute  $\min_{b \in B} \delta(r_0, b)$  for a given  $r_0 \in R$ . The other has to compute  $\min_{r \in R} \delta(r, b_0)$  for a given  $b_0 \in B$ . For the latter data structure we could use the dynamic nearest neighbour data structure by Chan [6], yielding polylogarithmic update and query times. However, for the former we need a dynamic weighted Voronoi diagram, and for this we only have the data structure developed by Agarwal, Efrat and Sharir [2]. A dynamic data structure for dynamic weighted Voronoi diagrams with updates and queries in polylogarithmic time readily would lead to  $\mathcal{O}(n \text{ polylog } n)$ .

### 3.3 Correctness and Complexity

Note that in the algorithm a point can only go from red to blue and from blue to dead. Dead points stay dead. We first prove two minor properties.

**Lemma 6.** *Once a point  $b^*$  is moved from  $B$  to  $D$ , it no longer has any edges to points in  $R$ .*

*Proof.* The move of  $b^*$  from  $B$  to  $D$  is a consequence of two facts: i)  $(b^*, r^*)$  is a BCP in  $B \times R$  — achieving the minimum of the expression

$$\min_{r \in R, b \in B} \delta(r, b) = \min_{r \in R, b \in B} \{\text{dist}[b] + \|r - b\|\},$$

and ii)  $\|b^* - r^*\| > 1$ . Therefore,  $\forall r \in R: \|b^* - r\| > 1$ .  $\square$

**Lemma 7.**  *$G(P)$  is not connected if and only if there is a moment in the algorithm when it holds that  $R \neq \emptyset$  and  $B = \emptyset$ .*

*Proof.* ( $\Rightarrow$ ): If  $G(P)$  is not connected, there is a point  $r$  that begins in  $R$  and is not reachable from  $s$ . It never leaves  $R$ , so  $R$  stays nonempty throughout.  $B$  gets emptied to  $D$  once the data structure starts returning only BCPs that do not form an edge in  $G(P)$ .

( $\Leftarrow$ ): By Lemma 6 the dead points do not have any edges to red points. If there are no blue points then  $G(P)$  is not connected.  $\square$

**Lemma 8.** *The algorithm WEIGHTEDSHORTESTPATHS correctly computes the shortest distances from the source and the parents of points in a SSSP tree.*

*Proof.* As in Dijkstra's algorithm, we find the vertex  $r^*$  minimizing the expression  $\text{dist}[b] + \|b - r\|$  over all vertices  $b \in B \cup D$  and  $r \in R$  with  $\|b - r\| \leq 1$ , and update the information of  $r^*$  accordingly. Thus the correctness follows from the correctness of Dijkstra's algorithm.  $\square$

**Lemma 9.** *The algorithm WEIGHTEDSHORTESTPATHS runs in  $\mathcal{O}(n^{1+\varepsilon})$  time and space, for an arbitrary constant  $\varepsilon > 0$ .*

*Proof.* The outer while loop runs at most  $2n - 2$  times, as in each iteration either a blue point is deleted and placed among the dead, or a red point becomes blue. If  $G(P)$  is not connected, the loop terminates even earlier by Lemma 7. In each iteration, either we finish because  $B = \emptyset$ , or we spend  $\mathcal{O}(1)$  time plus the time to make  $\mathcal{O}(1)$  operations in the BCP dynamic DS. Since by Theorem 5 each operation in the BCP dynamic DS takes  $\mathcal{O}(n^\varepsilon)$  amortized time, the result follows.  $\square$

**Theorem 10.** *Let  $P$  be a set of  $n$  points in the plane,  $s \in P$ , and  $\varepsilon > 0$  an arbitrary constant. The algorithm `WEIGHTEDSHORTESTPATHS`( $P, s$ ) returns the correct distances from the source in the graph  $G(P)$  in  $\mathcal{O}(n^{1+\varepsilon})$  time.*

*Proof.* By Lemmata 8 and 9. □

## 4 Conclusions

We have given algorithms to compute shortest paths in unit disk graphs in near-linear time. For the unweighted case it is easy to show that our algorithm is asymptotically optimal in the algebraic decision tree. A simple reduction from the problem of finding the maximum gap in a set of numbers shows that deciding if  $G(P)$  is connected requires  $\Omega(n \log n)$  time. As discussed in the text, a better data structure to dynamically maintain the bichromatic closest pair would readily imply an improvement in our time bounds for the weighted case.

A generalization of the graph  $G(P)$  is the graph  $G_{\leq t}(P)$ , where two points are connected whenever their distance is at most  $t$ . Thus  $G(P)$  is  $G_{\leq 1}(P)$ . Two natural extensions of our results come to our mind.

- Can we compute efficiently a compact representation of the distances in all the graphs  $G_{\leq t}(P)$ ?
- Can we find, for a given  $u, v, k$ , the minimum  $t$  such that in  $G_{\leq t}(P)$  the distance from  $u$  to  $v$  is at most  $k$ ? This problem can be solved in roughly  $\mathcal{O}(n^{4/3})$  using [1, 14] to guide a binary search in the interpoint distances. Can it be solved in near-linear time?

## Acknowledgments

We would like to thank Timothy Chan and David Eppstein for several useful comments. In particular, we are indebted to Timothy Chan for pointing out the work of Roditty and Segal [15].

## References

- [1] P. K. Agarwal, B. Aronov, M. Sharir, and S. Suri. Selecting Distances in the Plane. *Algorithmica*, 9(5):495–514, 1993.
- [2] P. K. Agarwal, A. Efrat, and M. Sharir. Vertical Decomposition of Shallow Levels in 3-Dimensional Arrangements and Its Applications. *SIAM J. Comput.*, 29(3):912–953, 1999.
- [3] P. K. Agarwal, M. H. Overmars, and M. Sharir. Computing Maximally Separated Sets in the Plane. *SIAM J. Comput.*, 36(3):815–834, 2006.

- [4] P. Bose, A. Maheshwari, G. Narasimhan, M. H. M. Smid, and N. Zeh. Approximating geometric bottleneck shortest paths. *Comput. Geom.*, 29(3):233–249, 2004.
- [5] P. B. Callahan and S. R. Kosaraju. A Decomposition of Multidimensional Point Sets with Applications to k-Nearest-Neighbors and n-Body Potential Fields. *J. ACM*, 42(1):67–90, 1995.
- [6] T. M. Chan. A dynamic data structure for 3-D convex hulls and 2-D nearest neighbor queries. *J. ACM*, 57(3), 2010.
- [7] T. M. Chan and A. Efrat. Fly Cheaply: On the Minimum Fuel Consumption Problem. *J. Algorithms*, 41(2):330–337, 2001.
- [8] B. N. Clark, C. J. Colbourn, and D. S. Johnson. Unit disk graphs. *Discrete Mathematics*, 86(13):165 – 177, 1990.
- [9] A. Efrat, A. Itai, and M. J. Katz. Geometry Helps in Bottleneck Matching and Related Problems. *Algorithmica*, 31(1):1–28, 2001.
- [10] D. Eppstein. Dynamic Euclidean Minimum Spanning Trees and Extrema of Binary Functions. *Discrete & Computational Geometry*, 13:111–122, 1995.
- [11] D. Eppstein. Testing bipartiteness of geometric intersection graphs. *ACM Transactions on Algorithms*, 5(2), 2009.
- [12] J. Gao and L. Zhang. Well-Separated Pair Decomposition for the Unit-Disk Graph Metric and Its Applications. *SIAM J. Comput.*, 35(1):151–169, 2005.
- [13] D. S. Hochbaum and W. Maass. Approximation Schemes for Covering and Packing Problems in Image Processing and VLSI. *J. ACM*, 32(1):130–136, 1985.
- [14] M. J. Katz and M. Sharir. An Expander-Based Approach to Geometric Optimization. *SIAM J. Comput.*, 26(5):1384–1408, 1997.
- [15] L. Roditty and M. Segal. On Bounded Leg Shortest Paths Problems. *Algorithmica*, 59(4):583–600, 2011.
- [16] P. M. Vaidya. Geometry Helps in Matching. *SIAM J. Comput.*, 18(6):1201–1225, 1989.